
FNIISC: Fault-Tolerant Networking through Intrusion Identification & Secure Compartments

Allison Mankin, Dan Massey, USC/ISI

Chien-Lung Wu, S. Felix Wu, Xiaoliang Zhao, UC Davis/NCSU

Dan Pei, Lan Wang, Lixia Zhang, UCLA

Randy Bush, AT&T (Project Consultant)

DARPA FTN PI Meeting, January 16, 2002

Building a More Resilient Internet

- Goal: build against attacks N times worse than 9/11 or Code Red
- How:
 - Identify fundamental pieces in the infrastructure
 - The current project focusing on the routing infrastructure
 - Assess how well each of them currently can resist faults/attacks
 - Build stronger and more fences to protect them
- Two of our recent results:
 - BGP: assessment of how well it stands against network attacks and failures now, what works and what to be improved
 - DNS: protecting DNS service from route hijacking

BGP Resilience during Code Red Attack

- Renesys report to NANOG, “*Global BGP routing instability during the Code Red attacks*”, showed
 - the correlation between the large attack traffic spikes due to the attack and BGP routing message spikes on 9/18/01
 - evidence of possibly large scale BGP *route* changes?
- Exactly how well/poorly did BGP actually behave as a routing protocol, and *why*?
 - Is BGP indeed in trouble facing virus attacks?
 - What insight about the routing protocol design can be inferred from the collected BGP data on 9/18?

Data and Methodology

- BGP update messages collected at RIPE NCC from 9/10/01 to 9/30/01


The same data set as collected by Renesys report

- 3 US peers: AT&T, Verio, Global Crossing
- 8 peers in Europe
- 1 in Japan

- Methodology


- ***Categorize*** BGP advertisements
- Infer the causes of each class of BGP advertisements

Prefix Announcements (15-min bins)



BGP Message Classification (1)

Prefix Updates in Each Class (All Peers)




What Our Analysis Shows

- A *substantial* percentage of the BGP messages during the worm attack were not about route changes
 - BGP initial table exchanges: 40.2% on 9/18/2001
 - Duplicate advertisements: 5% on 9/18/2001
- BGP updates that may indicate route changes:
 - Implicit withdraws: 37.6% on 9/18
- BGP updates that indicate route changes:
 - New Announcements: 8.8% on 9/18
 - Withdraws: 8.3% on 9/18

(roughly the same as during normal days)

View the data graphically ...

BGP Advertisements on 9/18/2001




A Closer Look at the Changes

- (40.2%) BGP table exchanges ← BGP session restarts
- (37.6%) Implicit withdraws
 - slow convergence
 - topology change
 - About 25% have unchanged ASPATH attribute
 - Most of them wouldn't be propagated by the receiver (e.g. changes in MED attribute)
 - Possible causes: internal network dynamics
- (~17%) Explicit route announcement and withdraws
 - Reachability and/or route changes
- (~5%) Duplicate announcements

How to Infer the Causes?

BGP Session Restarts → BGP table exchanges (40.2% of total BGP announcements on 9/18)




- All the BGP sessions restarted during the worm attack period (total 32 restarts)
- Synchronized session restarts on other days too
- *Largely an artifact of the monitoring point*

- Cern peering with NCC: 550K announcements due to session restart
- Cern peering with RRC04: 0

Implicit Withdraws (37.6% on 9/18)


Type 1: **same** ASPATH as the previous announcement (25%)

Type 2: otherwise





Type 1 Implicit Withdraw

Two US peers have high percentage of Type 1 Implicit Withdraws.



Type 2 Implicit Withdraw (28.9% on 9/18)

- The European peers have more Type 2 Implicit Withdraws than US peers.
- Causes: largely slow convergence?
 - more quantitative analysis coming.



One sample evidence of slow convergence

09/18/2001 14:04:23 A S3549 originated prefix 66.133.177.0/24


09/18/2001 14:04:37 A S1103 announced aspath 1103 3549

09/18/2001 14:05:10 A S3549 withdrew 66.133.177.0/24

09/18/2001 14:05:36 A S1103 announced aspath 1103 8297 6453 3549


09/18/2001 14:06:34 A S1103 announced aspath 1103 8297 6453 1239 3549

09/18/2001 14:07:02 A S1103 sent withdrawal to 66.133.177.0/24



How to Infer the Causes (2)

Duplicate Advertisements: One extreme Example: AT&T



About 30% of the total advertisements from AT&T were duplicate

What Insights Does This Give Us?

- BGP as a routing protocol stood well during the worm attack
 - It also stood up well during other major topological incidents, such as cable cuts, Baltimore tunnel fire, even 9/11 event.
- BGP design needs improvement for unforeseen future faults/attacks
 - BGP peering must work well not just on good days, but behave well even on rainy days
 - BGP should keep local changes *local* in order to be a more resilient *global* routing protocol
 - BGP fast convergence solution (that we reported in previous PI meeting) should be deployed to remove the "amplifier" effect of the slow route convergence under stressful conditions
- BGP implementation tradeoffs must be made in view of the system performance as a whole

Other Lessons Learned

- Be careful of measurement artifacts !
 - E.g. the impact of the sampling point: monitoring sites (mutli-hop eBGP) is different from the table exchanged used by actual point-to-point peers (direct exchange between adjacent links).
- Be careful of the distinction between the properties of a protocol and the behavior due to a particular implementation and/or configuration !
 - E.g. high duplicate updates from one service provider, uncontrolled flapping from another

What's interesting vs what's the challenge

- Thanks to the community effort, we do have some routing measurement data to look at now (provided by RIPE, Oregon Route Views, etc.)
 - One can generate lots interesting graphs
- Raw data alone does not necessarily tell *what is going on*, let alone *why*
- It is a great challenge to correctly *interpret* the data and *understand* the protocol in action
 - Strip off monitoring artifacts.
 - Strip off localized changes and errors.
 - Understand the dynamics of what the data means for the protocol ***in action***

Other Accomplishments

- Developed formal methods as tools which help reduce ambiguities in the BGP specification
- Evaluation of the MOAS solution design
 - simulation results show that this simple solution can effectively detect false routing announcements even in cases of multiple routers being compromised;
 - a partial deployment can substantially reduce the impact of false routing announcements
- Intention-driven itrace
 - FRiTrace package available
- Talks and Publications
 - Presentation at NANOG, October 2001
 - Submitted two IETF drafts (MOAS validation, itrace)
 - ICCCN'01 paper
 - "An Analysis of BGP Multiple Origin AS (MOAS) Conflicts" SIGCOMM Measurement workshop
 - Detection of Invalid Routing Announcements in the Internet (submitted to DSN 2002)
 - "On Fast BGP Convergence", to be presented at INFOCOM'02